arXiv:1410.7092v1 [cs.DS] 26 Oct 2014

Scheduling with Gaps: New Models and Algorithms

Marek Chrobak* Mordecai Golin† Tak-Wah Lam‡ Dorian Nogneng§

Abstract

We consider scheduling problems for unit jobs with release times, where the number or size of the gaps in the schedule is taken into consideration, either in the objective function or as a constraint. Except for a few papers on energy minimization, there is no work in the scheduling literature that uses performance metrics depending on the gap structure of a schedule. One of our objectives is to initiate the study of such scheduling problems with gaps. We show that such problems often lead to interesting algorithmic problems, with connections to other areas of algorithmics.

We focus on the model with unit jobs. First we examine scheduling problems with deadlines, where we consider variants of minimum-gap scheduling, including maximizing throughput with a budget for gaps or minimizing the number of gaps with a throughput requirement. We then turn to other objective functions. For example, in some scenarios, gaps in a schedule may be actually desirable, leading to the problem of maximizing the number of gaps. Other versions we study include minimizing maximum gap or maximizing minimum gap. The second part of the paper examines the model without deadlines, where we focus on the tradeoff between the number of gaps and the total or maximum flow time.

For all these problems we provide polynomial time algorithms, with running times ranging from $O(n \log n)$ for some problems, to $O(n^7)$ for other. The solutions involve a spectrum of algorithmic techniques, including different dynamic programming formulations, speed-up techniques based on searching Monge arrays, searching $X + Y$ matrices, or implicit binary search.

Throughout the paper, we also draw a connection between our scheduling problems and their continuous analogues, namely hitting set problems for intervals of real numbers. As it turns out, for some problems, the continuous variants provide insights leading to more efficient algorithms for the corresponding discrete versions, while in other problems completely new techniques are needed to solve the discrete version.

1 Introduction

We consider scheduling of unit jobs, with given release times, where the number or size of the gaps in the schedule is taken into consideration, either in the objective function or as a constraint. The problem of minimizing the number of gaps arises naturally in minimum-energy scheduling with deadlines, because in some scenarios the energy consumption associated with a schedule is a linear function of the number of gaps. To our knowledge, this is the only scheduling model in the literature that considers gaps in the schedule as a performance measure.

*Department of Computer Science, University of California at Riverside, USA. Research supported by NSF grants CCF-0729071 and CCF-1217314.

†Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Hong Kong.

‡Department of Computer Science, University of Hong Kong, Hong Kong.

§LIX, École Polytechnique, France.

As we show, however, one can formulate a number of other natural, but not yet studied variants of gap scheduling problems. Some of these problems can be solved using dynamic-programming techniques resembling those used for minimizing the number of gaps. Other require new techniques, giving rise to new and interesting algorithmic problems.

Throughout the paper, we focus on the model with unit-length jobs. The first type of problems we study involve unit jobs with release times and deadlines. In this category, we study the following problems:

- First, in Section 3, we briefly discuss minimizing the number of gaps for jobs with deadlines (under the assumption that the instance is feasible). An $O(n^4)$-time algorithm for this problem is known [3]. The main purpose of this section is to introduce the terminology and notation needed for subsequent sections, but we also include an alternative $O(n^4)$-time algorithm for this problem that is significantly simpler than the one in [3].

- We then address various extensions of the gap minimization problem. In Section 4, we study maximizing throughput (the number or total weight of executed jobs) with a budget γ for the number of gaps. We give an $O(\gamma n^6)$-time algorithm for this problem.

- In Section 5 we study the variant where we need to minimize the number of gaps under a throughput requirement, namely where either the number of jobs or their total weight must meet a specified threshold. We show that this problem can be solved in time $O(n^7)$.

- In the problems discussed earlier, the underlying assumption was that it is desirable to have as few gaps as possible. However, in certain applications gaps in a schedule may be actually desirable, which motivates the gap scheduling model where we wish to *maximize* the number of gaps. We study this problem in Section 6, and we provide an algorithm that computes an optimal schedule in time $O(n^6)$.

- Instead of the total number of gaps, the *size* of gaps may be a useful attribute of a schedule. In Section 7 we study the problem where, assuming that the given instance is feasible, we want to compute a schedule for which the maximum gap size is minimized. We give an $O(n^2 \log n)$-time algorithm for this problem.

We also consider scheduling problems where jobs have no deadlines. Now all jobs need to be scheduled. In this model we can of course schedule all jobs in one block, without gaps, but then some jobs may need to wait a long time for execution. To avoid this, we will also take into account the flow time measure, where the flow of a job is the time elapsed between its release and completion times, and we will attempt to minimize either the maximum flow or the total flow of jobs. We address three problems in this category:

- Minimizing total flow time with a budget γ for the number of gaps (Section 8). As we show, this problem can be solved in time $O(\gamma n + n \log n)$, by exploiting the Monge property of the dynamic programming arrays. The running time is in fact $O(\gamma n)$ if the jobs are given in sorted order of release times.

- Minimizing the number of gaps with a budget for total flow (Section 9). We show that this problem can be solved in time $O(n^2 \log n)$.

- Minimizing the number of gaps with a bound on the maximum flow time (Section 10). We show that this problem can be solved in time $O(n \log n)$, or even $O(n)$ if the jobs are already sorted in order of increasing release times.

- Minimizing maximum flow time with a budget γ for the number of gaps (Section 11). For this problem we give an algorithm with running time $O(n \log n)$.

Summarizing, for all these problems we provide polynomial time algorithms, with running times ranging from $O(n \log n)$ for some problems, to $O(n^7)$ for other. Interestingly, the solutions involve a spectrum of algorithmic techniques, including different dynamic programming formulations, and speed-up techniques based on searching Monge arrays, searching $X + Y$ matrices, and implicit binary search.

As another theme throughout the paper, we draw a connection between our scheduling problems and their continuous analogues, which are variants of the hitting set problems for intervals of real numbers, where jobs are represented by points on the real line. For example, the continuous version of the minimum-gap scheduling problem is equivalent to computing a hitting set of minimum cardinality. As it turns out, for some problems, the continuous variants provide insights leading to more efficient algorithms for the corresponding discrete versions, while in other problems completely new techniques are needed to solve the discrete version.

2 Preliminaries

The time is assumed to be discrete, divided into unit time intervals $[t, t+1)$, for $t = 1, 2, ...$, that we call *slots*. We will number these consecutive slots $0, 1, ...$, and we will refer to $[t, t+1)$ simply as *time slot* t, or occasionally even as *time* t. By $\mathcal{J}$ we will denote the instance, consisting of a set of unit-length jobs numbered $1, 2, ..., n$, each job j with a given integer release time r_j. This r_j denotes the first slot where j can be executed.

A *schedule* S of $\mathcal{J}$ is defined by an assignment of jobs to time slots such that (i) if a job j is assigned to a slot t then $t \geq r_j$, and (ii) no two jobs are assigned to the same slot. In most scheduling problems we consider we assume that all jobs are scheduled (assigned to slots). In problems that involve throughput, we will consider partial schedules, where only some of the jobs may be scheduled.

For a given schedule S, time slots where jobs are scheduled are called *busy*, while all other slots are called *idle*. A maximal finite time interval of idle slots in S is called a *gap* of S. A maximal time interval of busy slots is called a *block* of S. Of course, the number of blocks in S is always one more than the number of gaps.

Jobs with deadlines and feasibility. In some of the scheduling problems, the jobs in $\mathcal{J}$ will also have specified deadlines. The *deadline* of job j is denoted d_j, is assumed to be integer, and it is the last slot where j can be executed. (Thus it may happen that $d_j = r_j$, in which case j can only be executed in one slot.)

For instances with deadlines, we can restrict our attention to schedules S that satisfy the *earliest-deadline property (EDP)*: at any time t, either S is idle at t or it schedules a pending job with the earliest deadline. Using the standard exchange argument, any schedule can be converted into one that satisfies the earliest-deadline property and has the same set of busy slots. Without loss of generality, we can make the following assumptions about $\mathcal{J}$: (i) $r_j \leq d_j$ for each j, and (ii)

all jobs are ordered according to deadlines, that is $d_1 \leq \ldots \leq d_n$. Further, we can assume that (iii) all release times are distinct and all deadlines are distinct. The validity of assumption (iii) can be justified as follows. If two release times are equal, say when $r_i = r_j$ for $i \neq j$, and $d_i \leq d_j$, then we can modify the instance by setting $r_j = r_j + 1$. Symmetrically, if $d_i = d_j$ and $r_i \leq r_j$ then we can set $d_i = d_i - 1$. If this change produces a job with $d_i < r_i$, then job i cannot of course be scheduled. For problems where the feasibility is a requirement, we can then report that the instance is not feasible. For other problems, we can remove this job i from the instance altogether. It is easy to show, by a simple exchange argument, that the new schedule has the same throughput as the original one, so the correctness of our algorithms for the scheduling models with deadlines is not affected. These modifications can be done in time $O(n \log n)$, so they have no impact on the overall running time of our algorithms.

Instances without deadlines. For schedules involving objective functions other than throughput, we can assume that the jobs are ordered according to non-decreasing release times. For the total-flow objective function we can assume that all release times are different. This is because, although modifying the release times as above may change the total flow value, this change will be uniform for all schedules, so the optimality will not be affected. The maximum flow values though may change non-uniformly, so we will not be using this assumption in the Sections 10 and 11, where maximum flow of jobs is considered.

Shifting blocks. To improve the running time, some of our algorithms use assumptions about possible locations of the blocks in an optimal schedule. The general idea is that each block can be shifted, without affecting the objective function, to a location where it will contain either a deadline or a release time. The following lemma (that is implicit in [3]) is useful for this purpose. We formulate the lemma for leftward shifts; an analogous lemma can be formulated for rightward shifts and for deadlines instead of release times.

Lemma 1 *Assume that all jobs on input have different release times. Let $B = [u, v]$ be a block in a schedule such that the job scheduled at v has release time strictly before v. Then B can be shifted leftward by one slot, in the sense that the jobs in B can be scheduled in the interval $[u - 1, v - 1]$.*

Proof: We construct a sequence of job indices $i_1, i_2, ..., i_q$ such that i_1 is the job scheduled at v, each job i_b, $b = 2, 3, ..., q$ is scheduled in B at the release time $r_{i_{b-1}}$ of the previous job in the sequence, and $r_{i_q} < u$. This is quite simple: We let i_1 to be the job scheduled at v. Suppose that we already have chosen jobs $i_1, i_2, ..., i_c$, with each i_b scheduled at $r_{i_{b-1}}$, for $b = 2, 3, ..., c$. Since all release times are different, i_c must be released strictly before $r_{i_{c-1}}$. If $r_{i_c} < u$, we let $q = c$ and we are done. Otherwise, we have $u \leq r_{i_c} < r_{i_{c-1}}$. We then take i_{c+1} to be the job scheduled at r_{i_c}. By repeating this process, we obtain the desired sequence.

Given the jobs $i_1, i_2, ..., i_q$ from the previous paragraph, we can modify the schedule by scheduling i_q at time $u - 1$, and scheduling each i_b, $b = 1, 2, ..., q - 1$ at r_{i_b}. This will result in shifting B to the left by one slot, proving the lemma. □

Interval hitting. For some of these scheduling problems, it is useful to consider the "continuous" analogue of the problem. These continuous variants are obtained by assuming that all release times and deadlines are spread very far apart; thus in the limit we can think of jobs as having length 0. Each r_j and d_j (if deadlines are in the instance) is a point in time, and to "schedule" j we choose a

point in the interval $[r_j, d_j]$. Two jobs scheduled right next to each other end up being on the same point. This problem is then equivalent to computing hitting sets for a given collection of intervals on the real line, with some conditions involving gaps between the points in the hitting set.

More formally, in the hitting-set problem we are given a collection of intervals $I_j = [r_j, d_j]$ on the real line, where r_j, d_j are real numbers. (For scheduling problems without deadlines, we can let all d_j to be some sufficiently large number or ∞.) Our objective is to compute a set H of points such that $H \cap I_j \neq \emptyset$ for all j. This set H is called a *hitting set* of the intervals $I_1, I_2, ..., I_n$.

If H is a hitting set of $I_1, I_2, ..., I_n$, then for each j we can pick a *representative* $h_j \in H \cap I_j$. Sorting these representatives from left to right, $h_{i_1} \leq h_{i_2} \leq ... \leq h_{i_n}$, the non-zero differences between consecutive representatives, that is numbers $h_{i_{b+1}} - h_{i_b}$ are called *gaps* of H.

For each gap scheduling problem we can then consider the corresponding hitting-set problem. For example, we may want to compute a hitting set where the maximum gap in H is minimized.

These interval-hitting problems are conceptually easier to deal with than their discrete counterparts. As we show, some algorithms for interval-hitting problems extend to the corresponding gap scheduling problems, while for other the discrete cases require different techniques.

3 Minimizing the Number of Gaps

We first consider the minimum gap scheduling problem, where we have a collection of unit jobs with release times and deadlines. Assuming that the instance is feasible, the objective is to compute a schedule for all jobs that has a minimum number of gaps. This problem is closely related to minimum-energy scheduling in the power-down model, where a schedule specifies not only execution times for jobs but also at what times the processor can be turned off. The processor uses energy at rate L per time unit when the power is on, and it does not consume any energy when it is off. If the energy required to power-up the system is less than L then the energy minimization is equivalent to minimizing the number of gaps in the schedule. The problem was proposed by Irani and Pruhs, and it remained open for a few years. The first polynomial time algorithm, with running time $O(n^7)$, was given by Baptiste [2]. The time complexity was subsequently reduced to $O(n^4)$ in [3]. A greedy algorithm was analyzed in [4] and shown to have approximation ratio 2. Other variants of this problem have been studied, for example the multiprocessor case [5] or the case when the jobs have agreeable deadlines [1].

We include the discussion of the gap-minimization case mainly for the sake of completeness, and to introduce the reader to the dynamic-programming approach used to solve problems of this nature. However, we also include an alternative $O(n^4)$-time algorithm for this problem, which, although it does not improve the time complexity, is substantially simpler than the one in [3].

Continuous case. The continuous variant of minimum-gap scheduling is equivalent to computing a minimum-cardinality hitting set for a collection of intervals. It is well known (folklore) that this problem can be solved with a greedy algorithm in time $O(n \log n)$: Initialize $H = \emptyset$. Then, going from left to right, at each step locate the earliest-ending interval I_j not yet hit by the points in H and add d_j to H.

Scheduling. In the discrete case, different jobs need to be scheduled at different times, so the above greedy algorithm does not apply. No simple solutions to this problem are in fact known, both algorithms in [2, 3] are quite involved and are based on rather intricate dynamic programming

formulations. We now present an alternative algorithm for this problem. The fundamental approach we use is similar to that in [3], but the resulting algorithm and its analysis are substantially simpler.

Throughout this section we assume that $\mathcal{J}$ is feasible. As explained in Section 2, feasibility can be verified in time $O(n \log n)$. We can assume that jobs 1 and n are tight jobs that need to be executed, respectively, first and last, separated from the rest of the schedule, that is $r_1 = d_1 = \min_{j>1} r_j - 2 = 1$ and $d_n = r_n = \max_{j<n} d_j + 2$. We can always add two such jobs to the instance, as this increases the number of gaps by exactly two, independently of how other jobs are scheduled.

For $a, b, k \in \{1, 2, ..., n\}$, let $\mathcal{J}_{k,a,b}$ denote the sub-instance of $\mathcal{J}$ that consists of all jobs $j \in \{1, 2, ..., k\}$ that satisfy $r_a < r_j < r_b$. With $\mathcal{J}_{k,a,b}$ we will associate two functions

$\mathsf{G}_{k,a,b}$ = the minimum number of gaps in a schedule of $\mathcal{J}_{k,a,b}$ that schedules all jobs in the interval $[r_a + 1, r_b - 1]$, where we can count the initial gap starting at $r_a + 1$ (if there is one), but do not count the final gap, that is the one ending in $r_b - 1$ (if there is one).

$\mathsf{S}_{k,a,b}$ = the latest completion time $t \leq r_b - 1$ of a schedule with $\mathsf{G}_{k,a,b}$ gaps that satisfies the same restrictions as above. We refer to this value as the *stretch* of $\mathcal{J}_{k,a,b}$.

We wish to formulate a recurrence that would allow us to compute all values $\mathsf{G}_{k,a,b}$. The idea is this: If $k \notin \mathcal{J}_{k,a,b}$ then $\mathsf{G}_{k,a,b} = \mathsf{G}_{k-1,a,b}$. If $k \in \mathcal{J}_{k,a,b}$ then let k be scheduled at some time t. By the EDF property, if k is not the last job scheduled in $[r_a + 1, r_b - 1]$, and if there is a job c scheduled at time $t + 1$ then c cannot be available at time t. So $t = r_c - 1$. Furthermore, all jobs scheduled between r_a and $t - 1$ must be in $\mathcal{J}_{k-1,a,c}$ and all jobs scheduled between r_c and r_b are in $\mathcal{J}_{k-1,c,b}$, so we can compute the number of gaps in these intervals recursively. If there is no job scheduled at time $t+1$ then we can assume that k is scheduled as the last job in $[r_a + 1, r_b - 1]$. This is actually the case that creates complications for dynamic programming. In our algorithm, to deal with this case, we use the other function, $\mathsf{S}_{k,a,b}$.

Algorithm MinGaps**.** We process all instances $\mathcal{J}_{k,a,b}$ in order of increasing k, and for each k in order of increasing interval length, $r_b - r_a$, with $r_b - r_a > 0$. For each instance, we compute first $\mathsf{G}_{k,a,b}$ and then $\mathsf{S}_{k,a,b}$. We now give recurrence relations for these functions.

If $k = 0$, then we let $\mathsf{G}_{k,a,b} = 0$ and $\mathsf{S}_{k,a,b} = r_a$. If there is no release time between r_a and r_b then for all k, $\mathsf{G}_{k,a,b} = 0$ and $\mathsf{S}_{k,a,b} = r_a$. If $k \notin \mathcal{J}_{k,a,b}$ then $\mathsf{S}_{k,a,b} = \mathsf{S}_{k-1,a,b}$. Otherwise, $\mathsf{S}_{k,a,b} = \mathsf{S}_{k-1,a,b} + 1$.

For $|\mathcal{J}_{k,a,b}| \geq 1$, assume that the values of $\mathsf{G}_{k',a',b'}$ and $\mathsf{S}_{k',a',b'}$ for $k' < k$ are already computed. If $k \notin \mathcal{J}_{k,a,b}$ then $\mathsf{G}_{k,a,b} = \mathsf{G}_{k-1,a,b}$. Otherwise, we have $k \in \mathcal{J}_{k,a,b}$, in which case we let

$$\mathsf{G}_{k,a,b} = \min \left\{ \begin{array}{c} \min\limits_{\substack{c<k\,,\, r_k<r_c<r_b \\ \mathsf{S}_{k-1,a,c} \geq r_c - 2}} \{\mathsf{G}_{k-1,a,c} + \mathsf{G}_{k-1,c,b}\} \\ \mathsf{G}_{k-1,a,b} + \langle \text{if } \mathsf{S}_{k-1,a,b} + 1 < r_k \text{ then 1 else 0} \rangle \end{array} \right\}$$

To compute $\mathsf{S}_{k,a,b}$, we consider all choices, in the formula above, that give the minimum number of gaps $g = \mathsf{G}_{k,a,b}$ and we maximize the corresponding stretch values:

- In the top case, where k is scheduled right before r_c, the stretch value would be $\mathsf{S}_{k-1,c,b}$;
- In the bottom case, where k is scheduled at the end, the stretch value would be d_k if $\mathsf{S}_{k-1,a,b} + 1 < r_k$ (since k will create a new gap anyway), and $\mathsf{S}_{k-1,a,b} + 1$ otherwise.

The size of the dynamic programming tables $\mathsf{G}_{k,a,b}$ and $\mathsf{S}_{k,a,b}$ is $O(n^3)$. Computing each value takes time $O(n)$. Thus the overall running time of the algorithm is $O(n^4)$.

4 Maximizing Throughput with Budget for Gaps

In this section we consider a variant of gap scheduling where we want to maximize throughput, given a budget γ for the number of gaps. We first show that the continuous version of this problem can be solved in time $O(\gamma n^2)$. For the discrete case we give an algorithm with running time $O(\gamma n^6)$.

Continuous case. Formally, the problem is defined as follows. We are given a collection of intervals $I_j = [r_j, d_j]$, $j = 1, 2, ..., n$ and a positive integer γ. The objective is to compute a set H of at most γ points that hits the maximum number of intervals. As explained earlier, without loss of generality we only need to consider sets $H \subseteq \{d_1, d_2, ..., d_n\}$.

There is a fairly simple dynamic algorithm for this problem, that works as follows. Order the intervals according to the deadlines, that is $d_1 \leq d_2 \leq ... \leq d_n$. For $g = 1, 2, ..., \gamma$ and $b = 1, 2, ..., n$, let $\mathsf{T}_{b,g}$ be the maximum number of input intervals that are hit with at most g points from the set $\{d_1, d_2, ..., d_b\}$, assuming that one of these points is d_b. For all b, we first initialize $\mathsf{T}_{b,1}$ to be the number of intervals that contain d_b. Then, for $g \geq 2$ and all b, we can compute $T_{b,g}$ using the recurrence:

$$\mathsf{T}_{b,g} = \max_{a<b} \{\mathsf{T}_{a,g-1} + \Delta_{a,b}\},$$

where $\Delta_{a,b}$ is the number of intervals I_i such that $d_a < r_i \leq d_b \leq d_i$, namely the intervals that are hit by d_b but not d_a. The output value is $\max_b \mathsf{T}_{b,\gamma}$.

With a bit of care all values $\Delta_{a,b}$ can be pre-computed in time $O(n^2)$. To accomplish this, first sort all release times and deadlines. For each a, consider only intervals I_i with $r_i > d_a$. Then start with $x = r_a$, and keep incrementing x to the next release time or deadline, whichever is earlier, at each step updating the number of intervals hit by x, and recording these values for each $x = d_b$. This sweep costs time $O(n)$.

This gives us an algorithm with running time $O(\gamma n^2)$, because we have $O(\gamma n)$ values $\mathsf{T}_{b,g}$ to compute, each computation taking time $O(n)$.

Discrete case. For the discrete case, when we schedule unit jobs, a more intricate dynamic programming approach is needed. The fundamental approach is similar to that used in Algorithm MinGaps.

As before, denote by $\mathcal{J}$ the set of jobs on input. For each job k and times $u \leq v$, let let $\mathcal{K}_{k,u,v}$ denote the sub-instance of $\mathcal{J}$ that consists of all jobs $j \in \{1, 2, ..., k\}$ that satisfy $u \leq r_j \leq v$. (Note that this is a slightly different definition that in Section 3.) Define $\mathsf{T}_{k,u,v,g}$ to be the maximum number of jobs from $\mathcal{K}_{k,u,v}$ that can be scheduled in the interval $[u, v]$ with the number of gaps not exceeding g. Here, the initial and final gap (between u and the first job, and between the last job and u) are also counted.

The idea to compute $\mathsf{T}_{k,u,v,g}$ is this: Consider an optimal schedule for $\mathcal{K}_{k,u,v}$. If k is not scheduled, then $\mathsf{T}_{k,u,v,g} = \mathsf{T}_{k-1,u,v,g}$. So assume that k is scheduled, say at time t, where $u \leq t \leq v$. Letting h be the number of gaps before t, we have $\mathsf{T}_{k,u,v,g} = \mathsf{T}_{k-1,u,t-1,h} + \mathsf{T}_{k-1,t+1,v,g-h}$. If k is scheduled last, then either $d_k \geq v$ and k is scheduled at v, in which case $\mathsf{T}_{k,u,v,g} = \mathsf{T}_{k-1,u,v-1,g}$, or k is scheduled at some time $t' < v$ and $\mathsf{T}_{k,u,v,g} = \mathsf{T}_{k-1,u,t'-1,g-1}$.

Algorithm MaxThrpt**.** For all $k = 0, 1, ..., n$ and time slots u, v, where $v \geq u$, we process all instances $\mathcal{K}_{k,u,v}$ in order of increasing k, and for each k in order of increasing interval length, $v - u$. (We will explain shortly how to choose u and v to make the algorithm more efficient.) For each instance $\mathcal{K}_{k,u,v}$ and each gap budget g we compute the corresponding value $\mathsf{T}_{k,u,v,g}$.

First, if $\mathcal{K}_{k,u,v} = \emptyset$, we let $\mathsf{T}_{k,u,v,g} = 0$. Assume that $\mathcal{K}_{k,u,v} \neq \emptyset$. If $k \notin \mathcal{K}_{k,u,v}$ (which means that $r_k \notin [u, v]$) then $\mathsf{T}_{k,u,v,g} = \mathsf{T}_{k-1,u,v,g}$. Otherwise, we use the following recurrence:

$$\mathsf{T}_{k,u,v,g} = \max \left\{ \begin{array}{ll} \mathsf{T}_{k-1,u,v,g} & \\ \max\limits_{\substack{r_k \le t \le \min(d_k, v) \\ 0 \le h \le g}} \{\mathsf{T}_{k-1,u,t-1,h} + \mathsf{T}_{k-1,t+1,v,g-h}\} & \\ \max\limits_{r_k \le t' \le \min(d_k, v-1)} \left\{\mathsf{T}_{k-1,u,t'-1,g-1}\right\} & \\ \mathsf{T}_{k-1,u,v-1,g} & \text{if } d_k \ge v \end{array} \right\}$$

To achieve polynomial running time we need to show that only a small set of time parameters u, v and t needs to be considered. We claim that, without loss of generality, u and t are of the form $r_i - 1$ and v and t' are of the form $r_i + \sigma$, where $-n \le \sigma \le n$. The argument is as follows: We can assume that the optimal schedule satisfies the EDF property. So when k is scheduled in the second option in the maximum, then k is the only job pending at t. Further, we can assume that there is a job scheduled right after k, for otherwise we could reschedule k later, without increasing the number of gaps. This job, say r_c, will have to be scheduled at its release time r_c, which is equal $t + 1$. Thus t is of the form we claimed, which implies that u can also be assumed to be of this form. The claimed property of v and t' follows from the fact that we only need to consider schedules where each block contains a release time.

With the above restrictions on u, v and t, we have $O(n)$ choices for u and $O(n^2)$ choices for v. With $O(n)$ choices for k and γ choices for g, the size of the $\mathsf{T}_{k,u,v,g}$ table is $O(\gamma n^4)$. In the minimum, we either iterate over up to n choices for t and γ choices for h, or over up to n^2 choices for t'. So the overall running time is $O(\gamma n^6)$.

Weighted throughput. Both algorithms, for the continuous and discrete cases, easily extend to the model where jobs have non-negative weights and the objective is to maximize the weighted throughput. The running times remain the same.

5 Minimizing the Number of Gaps with Throughput Requirement

Suppose that now we want to minimize the number of gaps, under a throughput requirement, that is we want to find a schedule that schedules at least m jobs and minimizes the number of gaps. We assume that the maximum throughput is at least m; this condition can be verified in time $O(n \log n)$ by the earliest-deadline-first algorithm.

We can solve this problem, both the continuous and discrete version, by using the algorithms from the previous section. We explain the solution for the continuous variant; the solution of the discrete case can be obtained in a similar manner.

Recall that $\mathsf{T}_{b,g}$ was defined to be the maximum number of intervals that are hit with at most g points from the set $\{d_1, d_2, ..., d_b\}$, assuming that one of them is d_b. We can use these values to compute T_g, which is the maximum number of intervals hit with at most g points (without any restriction on the last point). Then, given our requirement m on the throughput, we compute the smallest g for which $\mathsf{T}_g \ge m$. This g is the output of the algorithm. The total running time will be $O(n^3)$.

An analogous approach will give us an algorithm for the discrete case with running time $O(n^7)$. In both cases, the algorithms extend to the model with weighted throughput without increasing the running time.

6 Maximizing the Number of Gaps

In the preceding sections we studied problems where we were looking for schedules with as few gaps as possible. However, in some applications, gaps in the schedule may actually be desirable. This arises, for example, when the input stream consists of high-priority jobs that are executed in pre-computed slots and low-priority jobs that are executed in slots not reserved by high-priority jobs. One such specific scenario appears in coordinating access to a channel through point coordination function (PCF) [7] where gaps (in our terminology) in the high-priority traffic are forced to allow low-priority traffic to access the channel.

Thus in this section we will examine the variant of gap scheduling where the objective is to create as many gaps as possible in the schedule. The continuous version of this problem is trivial: for any interval $I_j = [r_j, d_j]$ with $r_j = d_j$, we must of course choose $h_j = r_j$. Each interval $I_j = [r_j, d_j]$ with $r_j < d_j$ can be assigned a unique point $h_j \in I_j$. Thus in this section we will focus only on the discrete model.

Specifically, we are again given an instance with n unit jobs with release times and deadlines, and we assume that the instance is feasible, that is all jobs can be scheduled. The objective is to find a schedule that maximizes the number of gaps.

As before, we will assume that all jobs have different deadlines and different release times, and that they are ordered according to increasing deadlines. We can also assume that jobs 1 and n satisfy $r_1 = d_1 = \min_{j>1} r_j - 2$ and $d_n = r_n = \max_{j<n} d_j + 2$, that is, they are tight jobs executed at the beginning and end of the schedule, separated by gaps from other jobs.

For any job $k = 1, 2..., n$ and two time steps $u \leq v$ define $\mathcal{K}_{k,u,v}$ to be the sub-instance of $\mathcal{J}$ that consists of all jobs $j \in \{1, 2, ..., k\}$ that satisfy $u \leq r_j \leq v$. Define $\mathsf{D}_{k,u,v}$ to be the maximum number of gaps in a schedule of $\mathcal{K}_{k,u,v}$ in the interval $[u, v]$. In $\mathsf{D}_{k,u,v}$ we include the extremal gaps in the schedule, namely the gap between u and the first job and the gap between the last job and v.

Lemma 2 *For any sub-instance $\mathcal{K}_{k,u,v}$ there is a schedule S with $\mathsf{D}_{k,u,v}$ gaps in the interval $[u, v]$ that has the EDF property and satisfies the following two conditions:*

(i) For any job $j \in \mathcal{K}_{k,u,v}$, if j is scheduled at time S_j then all gaps in the interval $[r_j, S_j]$ have length at most 2 (including the gap between r_j and the first job).

(ii) For each block B of S, either all jobs in B are scheduled at their release times or, assuming that B does not start at u, the gap immediately to the left of B has length 1.

Proof: As explained in Section 2, each schedule can be converted into an EDF schedule without changing the set of busy slots, so we can assume that S has the EDF property.

It remains to show that we can modify any schedule S to have properties (i) and (ii). First, if some job j violates property (i), suppose that $[x, x']$ is a gap with $r_j \leq x < x + 2 \leq x' \leq S_j - 1$. We can then move j to time slot $x + 1$. Removing j from time slot S_j can decrease the number of gaps at most by 1 (if j was in a block by itself). Rescheduling j at time $x + 1$ will increase the number of gaps by 1. Thus overall the number of gaps cannot decrease.

If S has a block $B = [y, y']$ that violates property (ii), let l be the first job in B scheduled strictly after its release time r_l. Since all release times are different, we must have $r_l < y$. We can then move l to slot $y - 1$ and, since the gap that precedes B has length at least 2, the number of gaps will not decrease.

The two operations above convert the current schedule S into a new schedule S' whose set of busy slots is lexicographically smaller than that of S. The number of gaps in S' is the same or larger than the number of gaps in S. Thus this process must eventually end, producing a schedule that satisfies conditions (i) and (ii). $\square$

The fundamental idea of our algorithm is similar to the gap minimization algorithm from [2] and is based on dynamic programming, using Lemma 2 to achieve polynomial running time. We compute all values $\mathsf{D}_{k,u,v}$. We can assume that $k \in \mathcal{K}_{k,u,v}$, for otherwise $\mathsf{D}_{k,u,v} = \mathsf{D}_{k-1,u,v}$. Suppose that, in an optimal schedule, k is scheduled at some time $t \in [u, v]$. Then, by the EDF property, k cannot be a release time of any job $i \in \mathcal{K}_{k,u,v}$ other than k. Further, also using the EDF property, all jobs scheduled in $[u, t-1]$ belong to $\mathcal{K}_{k-1,u,t-1}$ and all jobs scheduled in $[t+1, v]$ belong to $\mathcal{K}_{k-1,t+1,v}$, thus giving us a partition of $\mathcal{K}_{k,u,v}$ into $\{k\}$ and two disjoint sub-instances $\mathcal{K}_{k-1,u,t-1}$ and $\mathcal{K}_{k-1,t+1,v}$, whose schedules must maximize the number of gaps in intervals $[u, t-1]$ and $[t+1, v]$, respectively. We thus conclude that $\mathsf{D}_{k,u,v} = \mathsf{D}_{k-1,u,t-1} + \mathsf{D}_{k-1,t+1,v}$. Since we do not know t a priori, we can maximize the expression on the right hand side over all choices of t. A complete description of the algorithm, including the initialization, is given below.

Algorithm MaxGaps**.** We consider all instances $\mathcal{K}_{k,u,v}$, where u and v are time slots such that $u, v \in [r_1, d_n]$ and $u \leq v+1$, and k is either a job, $k \in \{1, 2, ..., n\}$, or $k = 0$. We process these instances in order of increasing k and increasing difference $v-u$. For each instance $\mathcal{K}_{k,u,v}$, the value of $\mathsf{D}_{k,u,v}$ is computed as follows.

We first deal with boundary cases: If $u = v+1$ then $\mathsf{D}_{k,u,v} = 0$, and if $u \leq v$ and $k = 0$ then $\mathsf{D}_{k,u,v} = 1$.

So assume now that $u \leq v$ and $k \geq 1$. Then, if $k \notin \mathcal{K}_{k,u,v}$ we let $\mathsf{D}_{k,u,v} = \mathsf{D}_{k-1,u,v}$. Otherwise, we have $k \in \mathcal{K}_{k,u,v}$, in which case we compute $\mathsf{D}_{k,u,v}$ using the following recurrence:

$$\mathsf{D}_{k,u,v} = \max_{r_k \leq t \leq \min(v, d_k)} \{\mathsf{D}_{k-1,u,t-1} + \mathsf{D}_{k-1,t+1,v}\} \tag{1}$$

After all values are computed, the algorithm outputs D_{n,r_1,d_n}.

Implementation in time $O(n^5)$. The way the algorithm is stated, its running time cannot be expressed as a function of n. We now show how it can be implemented in polynomial time, by proving that we only need to consider a small set of candidates for u, v and t.

Let $V = \{r_j + c : j \in \{1, 2, ..., n\} \ \&\ c \in \{-1, 0, ..., 3n+1\}\}$. Clearly, $|V| = O(n^2)$. We claim that we only need to consider instances $\mathcal{K}_{k,u,v}$ where $u, v \in V$. This follows from Lemma 2(i), which implies that in the recurrence (1) for $\mathsf{D}_{k,u,v}$ we only need to consider slots t between r_k and $r_k + 3n$, inclusive. Thus, in the sub-instances $\mathcal{K}_{k-1,u,t-1}$ and $\mathcal{K}_{k-1,t+1,v}$ the new indices $v' = t-1$ and $u' = t+1$ will satisfy $v', u' \in \{r_k - 1, r_k - 2, ..., r_k + 3n + 1\} \subseteq V$. The final indices are r_1 and $d_n = r_n$, both in V, completing the proof of our claim. This gives us $O(n^5)$ instances $\mathcal{K}_{k,u,v}$ to consider, for $O(n^2)$ choices for each of u and v, and $O(n)$ choices for k. For each $\mathcal{K}_{k,u,v}$, to compute $\mathsf{D}_{k,u,v}$ we iterate only over $t = r_k, r_k + 1, ..., \min(v, d_k, r_k + 3n)$. Thus the overall running time is $O(n^6)$.

Next, we argue that this running time can be further improved to $O(n^5)$. The general idea is to show that, in essence, the recurrence (1) needs to be applied only to $O(n)$ values of u. We use a modified recurrence (1)

$$\mathsf{D}_{k,u,v} = \max_{r_k \leq t \leq \min(v, d_k, r_k+3n)} \left\{\mathsf{D}_{k-1,u,t-1} + \mathsf{D}_{k-1,\mu(t),v}\right\} \tag{2}$$

where $\mu(t)$ is computed as follows: If $t+1$ is a release time of a job in $\mathcal{K}_{k,u,v}$ then $\mu(t) = t+1$. If no job from $\mathcal{K}_{k,u,v}$ is released in $[t+1, v]$ then $\mu(t) = v+1$. Otherwise, $\mu(t) = r_i - 1$, where $i \in \mathcal{K}_{k,u,v}$ is the job with minimum release time r_i such that $r_i > t+1$. From the definition of $\mu(t)$, we have $\mathcal{K}_{k-1,t+1,v} = \mathcal{K}_{k-1,\mu(t),v}$ and, By Lemma 2, instance $\mathcal{K}_{k-1,\mu(t),v}$ is scheduled inside the interval $[\mu(t), v]$. So the algorithm remains correct if we restrict the range of u's to the set $U = \{r_j - c : j \in \{1, 2, ..., n\}\ \&\ c \in \{0, 1\}\}$, whose cardinality is $O(n)$. Then the total number of instances $\mathcal{K}_{k,u,v}$ to consider is $O(n^4)$, and the running time is $O(n^5)$.

7 Minimizing Maximum Gap

In earlier sections we focussed on the number of gaps in the schedule. For certain applications, the *size* of the gaps is also of interest. In this section we will study the problem where the objective is to minimize the maximum gap in the schedule. Such schedules will tend to produce many short gaps, which may be useful in applications discussed in Section 6, where a good schedule should leave some gaps between high-priority jobs, to allow other jobs to access the processor.

The general setting is as before. We have an instance $\mathcal{J}$ consisting of n unit jobs, where job j has release time r_j and deadline $d_j \geq r_j$. The objective is to compute a schedule of all jobs that minimizes the maximum gap size, or to report that no feasible schedule exists.

Interestingly, this problem is structurally different from these in the previous sections, because now, intuitively, a good schedule should spread the jobs more-or-less evenly in time. For example, if we have n jobs released at 0 and with deadline $D \gg n$, with a tight job at the beginning and a tight job at the end, we would then schedule them at times $i\frac{D}{n-1}$. In contrast, the algorithms in Sections 3–5 attempted to group the jobs into a small number of batches. Similar to the objective in Section 6, a schedule that minimizes the maximum gap size will typically create many gaps, but, as can be seen in Figure 1, these two objective functions will produce, in general, different schedules.

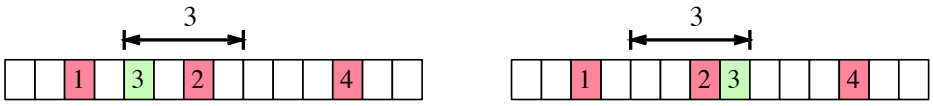

Figure 1: An instance with two schedules. (Red/dark shaded slots represent tight jobs.) The schedule on the left maximizes the number of gaps. The schedule on the right miminimizes the maximum gap. Both schedules are unique optimal solutions for their respective objective functions.

In this section we give an $O(n^2 \log n)$-time algorithm for computing schedules that minimize the maximum gap. We first give an algorithm for the continuous model, and then extend it to the discrete model.

7.1 The Continuous Case

The continuous analogue of our scheduling problem can be formulated as follows. The input consists of n intervals $I_1, I_2, ..., I_n$. As before, $I_j = [r_j, d_j]$, for each j. The objective is to compute a hitting set H for which the maximum gap between its points is minimized. Another way to think about

this problem is as computing a representative $h_j \in H \cap I_j$ for each interval I_j. Except for degenerate situations (two equal intervals of length 0), we can assume that all representatives are different, although we will not be using this property in our algorithm, and we treat H as a multiset.

Without losing generality, we can assume that $d_1 \leq d_2 \leq ... \leq d_n$. Further, we only need to be concerned with sets H that contain d_1, because if H contains any points before d_1 then we can replace them all by d_1 without increasing the maximum gap in H. Also, if $\max_i r_i \leq d_1$ then there is a singleton hitting set, $H = \{d_1\}$, in which case the maximum gap is vacuously equal 0. Thus we can also assume that $\max_i r_i > d_1$, so that we need at least two points in H.

Consider first the decision version: "Given $\lambda > 0$, is there a hitting set H for $I_1, I_2, ..., I_n$ in which all gaps are at most λ?" If λ has this property, we will call it *viable*. We will first give a greedy algorithm for this decision version and then later we will show how to use it to obtain an efficient algorithm for the minimization version.

Algorithm VIABLE(λ)**.** We initialize $h_1 = d_1$ and $U = \{2, 3, ..., n\}$. U is the set containing the indices of the intervals that still do not have representatives selected. We move from left to right, at each step assigning a representative to one interval in U, and we remove this interval from U.

Specifically, at each step, we proceed as follows. Let $z = \max(H) + \lambda$. If all $j \in U$ satisfy $r_j > z$, declare failure by returning false. Otherwise, choose $j \in U$ with $r_j \leq z$ that minimizes d_j and remove j from U. We now have two cases. If $d_j \leq z$, let $h_j = d_j$, and otherwise (that is, when $r_j \leq z < d_j$) let $h_j = z$. Then add h_j to H, and continue. If the process completes with $U = \emptyset$, return true. (The solution is $H = \{h_1, h_2, ..., h_n\}$.)

It is easy to prove that Algorithm VIABLE(λ) is correct. Let $H^* = \{h_1^*, h_2^*, ..., h_n^*\}$ be some solution with all gaps at most λ. We show that this solution can be converted into the one computed by the algorithm. For $j = 1$ we have $h_1^* = h_1 = d_1$ (as explained earlier). Consider the first step where Algorithm VIABLE chooses $h_j \neq h_j^*$. (If there is no such step, we are done.) By the choice of h_j in the algorithm, we have $h_j^* < h_j$. We can then replace h_j^* by h_j in H^*, without increasing the gap size to above λ. This way, we increase the number of steps of Algorithm VIABLE that produce the same representatives as those in H^*, and we can repeat the process.

We claim that the algorithm can be implemented in time $O(n \log n)$. Instead of U, the algorithm can maintain a set U' consisting only of those intervals j with $r_j \leq \max(H) + \lambda$ that that do not have yet a representative. Then choosing the j in the algorithm and removing it from U' takes time $O(\log n)$. When $\max(H)$ is incremented (after adding h_j), new intervals are inserted into U' in order of release times, each insertion taking time $O(\log n)$.

Now, the idea is to use the above procedure as an oracle and to do a binary search on λ's. To this end, we would like to have a small set of candidate values for the optimal λ. Let

$$\Lambda = \left\{ \frac{r_i - d_j}{k} : k \in \{1, 2, ..., n-1\}, i, j \in \{1, 2, ..., n\}, r_i > d_j \right\}.$$

We claim that Λ contains an optimal gap length λ^*. The argument is this. Consider some hitting set H with maximum gap λ^*, say $H = \{h_1, h_2, ..., h_n\}$. Let $h'_1 \leq h'_2 \leq ... \leq h'_n$ be H sorted in non-decreasing order. Choose some maximal (w.r.t. inclusion) consecutive sub-sequence $h'_a < h'_{a+1} < ... < h'_b$ with all gaps equal to λ^*, and suppose that h'_a is not a deadline. Then we can move h'_a by a little bit to the right without creating a gap longer than λ^*. Similarly, if h'_b is not a release time then we can apply a similar procedure to h'_b and shift it to the left. Each such operation reduces

the number of gaps of length λ^*. Since λ^* is optimal, eventually we must get stuck, meaning that we will find a sub-sequence like above with the first and last indices a and b that satisfy $h'_a = d_j$ and $h'_b = r_i$, for some i and j. Then we will have $\lambda^* = \frac{r_i - d_j}{b-a} \in \Lambda$.

The idea above immediately yields an $O(n^3 \log n)$-time algorithm. This algorithm first computes the set Λ, sorts it, and then finds the optimal λ through binary search in Λ. Note that the running time is dominated by sorting Λ.

Below we show that this running time can be improved to $O(n^2 \log n)$, by a more careful search in Λ that avoids constructing Λ explicitely.

Algorithm MinMaxGap**.** The $O(n^2 \log n)$-time algorithm is shown in Pseudocode 1. In this pseudo-code, to avoid multi-level nesting, we assume that the algorithm terminates if the **return** statement is reached. We now explain the steps in the algorithm and justify correctness and the running time.

Pseudocode 1 Algorithm MinMaxGap

1: **if** $\max_i r_i \leq d_1$ **then return** 0
2: $\Delta \leftarrow \{r_i - d_j : r_i > d_j, i, j \in 1, 2, ..., n\}$
3: sort Δ in non-decreasing order
4: **if** Viable($\frac{\min(\Delta)}{n-1}$) **then return** $\frac{\min(\Delta)}{n-1}$
5: $v \leftarrow \max\left\{u \in \Delta : \text{Viable}(\frac{v}{n-1}) = \mathsf{false}\right\}$
6: $w \leftarrow \min\{u \in \Delta : u > v\}$
7: **if** Viable(v) = false **then return** $\frac{w}{n-1}$
8: $k_0 \leftarrow \max\left\{k \in \{1, 2, ..., n-1\} : \text{Viable}(\frac{v}{k}) = \mathsf{true}\right\}$
9: $\Lambda' \leftarrow \left\{\frac{u}{\lceil k_0 u/v \rceil} : u \in \Delta \text{ and } \frac{v}{k_0+1} < u \leq v\right\} \cup \left\{\frac{w}{n-1}\right\}$
10: sort Λ' in non-decreasing order
11: **return** $\min\{\lambda \in \Lambda' : \text{Viable}(x) = \mathsf{true}\}$

First, if $\max_i r_i \leq d_1$ then there is a hitting set with all representatives on one point, so all gaps are 0. So suppose that $\max_i r_i > d_1$, that is any hitting set needs at least two points and the maximum gap is strictly positive. We then compute all positive values $r_i - d_j$, store them in a set Δ and sort them (Lines 3-4). This will take time $O(n^2 \log n)$.

If $\frac{\min(\Delta)}{n-1}$ is viable (which we check in Line 5), then this is the optimal value, since no hitting set can have all gaps smaller than $\frac{\min(\Delta)}{n-1} = \min(\Lambda)$. We can thus now assume that $\frac{\min(\Delta)}{n-1}$ is not viable.

Next, we compute the smallest $v \in \Delta$ for which $\frac{v}{n-1}$ is not viable. By the previous paragraph, such v exists. To this end, we can do binary search in the set $\left\{\frac{u}{n-1} : u \in \Delta\right\}$, at each step making calls to Viable() to determine whether the current split value is viable or not. The binary search will take time $O(n \log^2 n)$. We also let w to be the next value in Δ. (If there is no such value, let $w = \infty$.)

At this point we check whether v is viable. If it is not, it means that for all $u \leq v$, all values u/k, for $k = 1, 2, ..., n-1$, are not viable as well. Then the smallest viable value in Λ must be $\frac{w}{n-1}$,

so we output $\frac{w}{n-1}$ in Line 10. (Note that in this case w must exist, because the largest value in Δ must be viable.)

If v is viable, we compute the largest k_0 for which v/k_0 is viable. By the choice of v we have $k_0 < n-1$. We now also know that the optimal value for λ is in the interval $Q = (v/(k_0+1), v/k_0]$. Thus we only need to be considered with numbers in Λ that are in Q. We now claim that for any $u \in \Delta$, if $u \leq v$ then there is at most one integer $k_u \in \{1, ..., n-1\}$ for which $u/k_u \in Q$. This follows from simple calculation: We have

$$\frac{u}{k} \in Q \quad \text{implies} \quad \frac{v}{k_0+1} < \frac{u}{k} \leq \frac{v}{k_0} \quad \text{implies} \quad \frac{u}{v} \cdot k_0 \leq k \leq \frac{u}{v} \cdot k_0 + \frac{u}{v} \leq \frac{u}{v} \cdot k_0 + 1.$$

Thus the only candidate for k_u is $k_u = \lceil \frac{u}{v} \cdot k_0 \rceil$.

The above argument gives us that the only candidates for the optimal gap size we need to consider are all values u/k_u, for $u \in \Delta$ and $u \leq v$, plus the value w that we computed before. In Lines 12-13 we let Λ' be the set of these candidates and we sort them in non-decreasing order. Finally, we find the smallest viable $x \in \Lambda'$. This can be done in time $O(n \log^2 n)$ with binary search that calls Viable() for each split value at each step.

Overall, we obtain that Algorithm MinMaxGap correctly computes the optimal gap size and that its running time is $O(n^2 \log n)$. Note that this running time is dominated by sorting the sets of candidate gap sizes Δ and Λ'.

7.2 The Discrete Case

We now show that Algorithm MinMaxGap from the previous section can be adapted to the discrete case, namely to scheduling unit jobs.

Let $\mathcal{J}$ be an instance of unit job scheduling with release times and deadlines. We treat $\mathcal{J}$ as a collection of intervals $[r_j, d_j]$, $j = 1, 2, ..., n$, and run Algorithm MinMaxGap. This will produce the representatives $h_1, h_2, ..., h_n$ for the intervals in $\mathcal{J}$. Let λ be the maximum gap between these representatives. Since λ is an optimal gap for the continuous variant, it is a lower bound on the gap length for the discrete variant. It is thus enough to construct a schedule with gap at most $\bar{\lambda} = \lceil \lambda \rceil$.

To obtain a schedule, we simply schedule each job j at time $\bar{h}_j = \lceil h_j \rceil$. For each j we have $h_j \in [r_j, d_j]$, so $\bar{h}_j \in [r_j, d_j]$ as well, because all release times r_j and deadlines d_j are integral. So the schedule $\bar{h}_1, \bar{h}_2, ..., \bar{h}_n$ is feasible.

To estimate the gap size, consider some job i that is not last in the schedule and let j be the first job after i. We have $\bar{h}_j = \lceil h_j \rceil \leq \lceil h_i + \lambda \rceil \leq \lceil h_i \rceil + \lceil \lambda \rceil = \bar{h}_i + \bar{\lambda}$. Thus the gap in the schedule is at most $\bar{\lambda}$.

Summarizing, we obtain that, with the above modification, Algorithm MinMaxGap produces a schedule that minimizes the maximum gap and it runs in time $O(n^2 \log n)$.

8 Minimizing Total Flow Time with a Budget for Gaps

Unlike in earlier sections, we now consider jobs without deadlines and focus on the tradeoff between the number of gaps and delay. Formally, an instance $\mathcal{J}$ is given by a collection of n unit length jobs. For each job $j = 1, 2, ..., n$ we are given its release time r_j. If, in some schedule, job j is executed at time S_j then $F_j = S_j - r_j$ is called the *flow time of* j. We are also given a budget value γ. The objective is to compute a schedule for $\mathcal{J}$ in which the number of gaps is at most γ and the total flow time $\sum_j F_j$ is minimized.

Continuous case. The continuous variant of this problem is equivalent to the k-medians problem on a directed line: given points $r_1, r_2, ..., r_n$, find a set H of k points that minimizes the sum $\sum_{i=1}^{n} \min_{h \in H \wedge h \geq r_i}(h - r_i)$, where the ith term of the sum represents the distance between r_i and the first point in H after r_i. (Here, the value of k corresponds to $\gamma - 1$, the number of blocks.) This is a well-studied problem and it can be solved in time $O(kn)$ if the points are given in a sorted order [8].

Discrete case. The discrete case differs from its continuous analogue because the jobs executed in the same block do not occupy a single point. We show, however, that the techniques for computing k-medians can be adapted to minimum-flow scheduling with gaps, resulting in an algorithm with running time $O(n \log n + \gamma n)$.

Without loss of generality, we assume that all release times are different and ordered in increasing order, that is $r_1 < r_2 < ... < r_n$. Any instance can be modified to have this property in time $O(n \log n)$. As explained in Section 2, this modification changes the flow of all schedules uniformly, so the optimality is not affected. This is the only part of the algorithm that requires time $O(n \log n)$; the remaining part will run in time $O(\gamma n)$.

We first give a simple dynamic programming formulation with running time $O(\gamma n^2)$, and then show how to improve it to $O(\gamma n)$. Any schedule with at most γ gaps consists of at most $\gamma + 1$ blocks. To reduce the running time, we need to show that these blocks can only be located at a small number of possible places. For this, we will need the following lemma.

Lemma 3 *There is an optimal schedule with the following properties: (i) all jobs are scheduled in order of their release times, and (ii) the last job of each block is scheduled at its release time.*

Proof: Property (i) is trivial. Property (ii) follows by iteratively applying Lemma 1. □

Based on this lemma, each block consists of consecutive jobs, say $i, i+1, ..., j$, with the last job j scheduled at time r_j. The hth job of the block is scheduled at time $r_j - j + h$. So the contribution of this block to the total flow is

$$\begin{aligned} W_{i,j} = \sum_{h=i}^{j-1} F_j &= \sum_{h=i}^{j-1} (r_j - j + h - r_h) \\ &= (j-i) r_j - \binom{j-i+1}{2} - R_{j-1} + R_{i-1}, \end{aligned}$$

where $R_b = \sum_{a=1}^{b} r_a$, for each job b.

Algorithm MinTotFlow**.** Let $\mathsf{F}_{j,g}$ denote the minimum flow of a schedule for the sub-instance consisting of jobs $1, 2, ..., j$ with the number of gaps at most g. We initialize $\mathsf{F}_{0,g} = 0$ and $\mathsf{F}_{j,0} = W_{1,j}$ for $j > 0$. Then, for $j = 1, 2, ..., n$ and $g = 1, ..., \gamma$, we compute

$$\mathsf{F}_{j,g} = \min_{1 \leq i \leq j} \{\mathsf{F}_{i-1,g-1} + W_{i,j}\}.$$

To justify correctness, we need to explain why the above recurrence holds. Consider a schedule that realizes $\mathsf{F}_{j,g}$. From Lemma 3, since we are minimizing the total flow, we can assume that job j is scheduled at r_j. Let i be the first job of the last block. As we calculated earlier, the contribution

of this block to the total flow is $W_{i,j}$. The minimum flow of the remaining jobs $1, 2, ..., i-1$, is independent of how jobs $i, ..., j$ are scheduled, so (inductively) it is equal $\mathsf{F}_{i-1,g-1}$.

We now consider the running time. All values $W_{i,j}$ can be precomputed in time $O(n^2)$. We have at most γ choices for g and at most n choices for j, so there are $O(\gamma n)$ values $\mathsf{F}_{j,g}$ to compute. Computing each value takes time $O(n)$, for the total running time $O(\gamma n^2)$.

To improve the running time to $O(\gamma n)$, we show that the values $W_{i,j}$ satisfy the quadrangle inequality.

Lemma 4 *For all $1 \le i < j < n$, we have*

$$W_{i,j} + W_{i+1,j+1} \le W_{i,j+1} + W_{i+1,j}.$$

Proof: We compute $W_{i,j} - W_{i+1,j}$ and $W_{i+1,j+1} - W_{i,j+1}$ separately:

$$\begin{aligned} W_{i,j} - W_{i+1,j} &= \left[(j-i)r_j - \binom{j-i+1}{2} - R_{j-1} + R_{i-1}\right] \\ &\qquad - \left[(j-i-1)r_j - \binom{j-i}{2} - R_{j-1} + R_i\right] \\ &= r_j - j + i - r_i, \end{aligned}$$

and

$$\begin{aligned} W_{i+1,j+1} - W_{i,j+1} &= \left[(j-i)r_{j+1} - \binom{j-i+1}{2} - R_j + R_i\right] \\ &\qquad - \left[(j+1-i)r_{j+1} - \binom{j-i+2}{2} - R_j + R_{i-1}\right] \\ &= -r_{j+1} + j + 1 - i + r_i. \end{aligned}$$

Adding these inequalities, we get

$$W_{i,j} + W_{i+1,j+1} - W_{i,j+1} - W_{i+1,j} = r_j - r_{j+1} + 1 \le 0,$$

because $r_j < r_{j+1}$, due to our assumption that all release times are different. □

Lemma 4 implies that function $\mathsf{F}_{j,g}$ has the Monge property, which can be used to speed up the dynamic programming algorithm to $O(\gamma n)$, see [8]. Adding the time $O(n \log n)$ needed for preprocessing, the overall running time of Algorithm MinTotFlow is $O(n \log n + \gamma n)$.

9 Minimizing Number of Gaps with a Bound on Total Flow

An alternative way to formulate the tradeoff in the previous section would be to minimize the number of gaps, given a budget f for the total flow. This can be reduced to the previous problem by finding the smallest g for which there is a schedule with at most g gaps and flow at most f.

The solution is the same for both the continuous and discrete versions, so we focus on the discrete variant. Specifically, let $\hat{F}(g)$ be the minimum flow of a schedule that has at most g gaps. Then $\hat{F}(g)$ is monotonely decreasing as g increases. Using binary search, where at each step we use Algorithm MinTotFlow as an oracle, we can then find the smallest g for which $\hat{F}(g) \le f$. The resulting algorithm will have running time $O(n^2 \log n)$.

10 Minimizing Number of Gaps with a Bound on Maximum Flow

Now, instead of minimizing the *total* flow time, we consider the objective function equal to the *maximum* flow time, $F_{\max} = \max_j(S_j - r_j)$. At the same time, we would also like to minimize the number of gaps. This leads to two optimization problems, by placing a bound on one value and minimizing the other. In this section we consider the problem of minimizing the number of gaps when an upper bound on the flow of each job is given. For this problem, we give $O(n \log n)$-time algorithm.

Formally, we are given an instance consisting of n unit jobs with release times, and a threshold value f. The objective is to compute a schedule of these jobs in which each job's flow time is at most f and the number of gaps is minimized. As before, without loss of generality, we can assume that the jobs are sorted according to their release times, that is $r_1 \leq r_2 \leq ... \leq r_n$.

Continuous case. We start by giving an $O(n \log n)$-time algorithm for the continuous case. Here we are given a collection of n real numbers $r_1, r_2, ..., r_n$, and a number f, and we want to compute a set H of minimum cardinality such that $\min\{h \in H : h \geq r_i\} \leq r_i + f$, for all $i = 1, 2, ..., n$.

We show that this can be solved in time $O(n)$, assuming the input is sorted. Indeed, this is very simple, using a greedy algorithm that computes H in a single pass through the input. Specifically, initialize $H = \{r_1 + f\}$, and then in each step, choose i to be smallest index for which $r_i > \max(H)$ and add $r_i + f$ to H. A routine inductive argument shows that the computed set H has indeed minimum cardinality. If $r_1 \leq r_2 \leq ... \leq r_n$, then the algorithm is essentially a linear scan through this sequence, so its running time is $O(n)$. With sorting, the time will be $O(n \log n)$.

Discrete case. Next, we want to show that we can achieve the same running time for the discrete variant, where we schedule unit jobs. Our approach is to reduce the problem to the gap minimization problem for the version with deadlines. To this end, we define the deadline d_j for each job j as $d_j = r_j + f$. We now need to solve the gap minimization problem for jobs with deadlines, the problem we discussed earlier in Section 3, and for which we gave an algorithm with running time $O(n^4)$. However, we can do better than this. The instances we created satisfy the "agreeable deadline" property, which means that the ordering of the deadlines is the same as the ordering of release times. For such instances a minimum-gap schedule can be computed in time $O(n \log n)$ (unpublished folklore result, see [1]). This will thus give us an $O(n \log n)$-time algorithm for gap minimization with a bound on maximum flow.

In the remainder of this section we present an alternative $O(n \log n)$-time algorithm for this problem, which has the property that its running time is actually $O(n)$ if the jobs are already sorted in non-decreasing order of release times. This algorithm will be useful in the next section.

From now on we assume that the instance consists of n unit jobs numbered $1, 2, ..., n$, whose release times are ordered in non-decreasing order: $r_1 \leq r_2 \leq ... \leq r_n$. As we remarked earlier in Section 2, we cannot now assume that all jobs have different release times. In fact, the presence of jobs with equal release times causes the algorithm for the discrete case to be more involved than for the continuous case. It is easy to see (by a simple exchange argument) that, without loss of generality, we only need to consider schedules in which jobs are also scheduled in the same order $1, 2, ..., n$.

Algorithm MinGapMaxFlow**.** We will have two stages. In the first stage we produce a tentative schedule by greedily scheduling the jobs from left to right: At each time step t, if there is a pending job, we schedule the earliest-deadline pending job and increment t by 1. If there is no pending job, we set t to be the next release time. If there is a job in this schedule with flow larger than f, declare that there is no schedule with flow f and stop. If all jobs are scheduled, continue to the next stage. This stage clearly runs in time $O(n)$.

We now explain the second stage. In this stage we want adjust our tentative schedule, by shifting some jobs to the right to reduce the number of gaps, without exceeding the maximum flow restriction. The computation will consist of phases, where in each phase we construct one block of the schedule. Right before phase l, the schedule will satisfy the following invariant: all jobs are scheduled in blocks $B_1, B_2, ..., B_{l-1}$, from left to right, where $B_h = [u_h, v_h]$, $h = 1, 2, ..., l-1$. For convenience, we think of each block as a set of jobs and as an interval. These blocks will not change in the future. We will maintain the following invariant: (i) The jobs scheduled in these blocks are exactly the jobs released before v_{l-1} (inclusive); (ii) These jobs are scheduled in order of their release times; (iii) Each block has at least one job with flow f; and (iv) No job is scheduled at time $v_{l-1} + 1$ (in the tentative schedule).

Let i be the first job after B_{l-1}, scheduled at time S_i in the tentative schedule. From the invariant, $S_i > v_{l-1} + 1$. We start with $B_l = [u_l, v_l] = [S_i, S_i]$ consisting only of job i. Within this phase, in each step, we will either shift B_l to the right or add another job to it. Specifically, we do this. If there is a job scheduled at $v_l + 1$, we add this job to B_l and increment v_l. Suppose now that there is no job released at $v_l + 1$. If all jobs in B_l have flow time smaller than f then shift all jobs in B_l by 1 to the right and increment u_l and v_l. Otherwise, we end the phase.

We now claim that this algorithm can be implemented in time $O(n)$. For each block B_l we have a list of jobs scheduled in this block, in order of release times. Also, for each block we maintain the maximum flow time in this block, $F(B_l) = \max_{j \in B_l}(S_j - r_j)$. (However, we do not actually keep track of individual S_j's.) If we add a new job i to B_l then we update $F(B_l) \leftarrow \max(F(B_l), S_i - r_i)$, where S_i is the execution time of i in the tentative schedule. Instead of shifting B_l one by one, we compute the total shift after which either the maximum flow of B_l will become f or there will be a job scheduled right after B_l, whichever is smaller. Specifically, if i is the first job scheduled after B_l, at time S_i, then the shift value is $\delta = \min(S_i - v_l - 1, f - F(B_l))$, that is, both u_l and v_l are increased by δ. After this computation is complete, the schedule time of each job can be computed by adding its index within its block to this block's start time. With these modifications, the running time of the second stage will be $O(n)$.

11 Minimizing Maximum Flow with a Budget for Gaps

We now consider an alternative variant of the tradeoff between maximum flow and the number of gaps. This time, for a given collection of n unit jobs with release times $r_1, r_2, ..., r_n$, and a budget γ, we want to compute a schedule that minimizes the maximum flow time value $F_{\max}$. (Recall that $F_{\max} = \max_j(S_j - r_j)$, where S_j is the execution time of job j.)

Continuous case. In the continuous case, $r_1, r_2, ..., r_n$ are points on the real line, and we want to compute a set H of at most γ points that minimizes $F_{\max}(H) = \max_i \min_{x \in H, x \geq r_i} |x - r_i|$. (This is a special case of the k-center problem, when the underlying space is the directed line.) We can assume that $r_1 \leq r_2 \leq ... \leq r_n$. Then the idea of the algorithm is to do binary search for

the optimal value f^* of $F_{\max}(H)$, at each step of the binary search using the algorithm from the previous section as an oracle.

For binary search, we need a small set of candidate values for f^*. If H is an optimal set of points, then, without loss of generality, we can assume that H contains only release times, since any other point in H can be shifted left until it reaches a release time. Thus we only need to consider the set Φ of all values of the form $r_j - r_i$ for $j > i$. Since $|\Phi| = O(n^2)$ and we need to sort Φ before doing binary search, we would obtain an $O(n^2 \log n)$-time algorithm.

Fortunately, we do not need to construct Φ explicitly. Observe that the elements of Φ can be thought of as forming an $X + Y$ matrix with sorted rows and columns, where X is the vector of release times and $Y = -X$. We can thus use the $O(n)$-time selection algorithm for $X + Y$ matrices [6] to speed up computation. Specifically, at each step we will have two indices p, q, with $1 \leq p \leq q \leq n(n-1)/2$, and we will know that the optimal value of f^* is between the pth and qth smallest values in Φ, inclusive. We let $l = \lfloor (p+q)/2 \rfloor$ and we use the algorithm from [6] to find the lth smallest element in Φ, say f. We now determine whether $f^* \leq f$ by applying the $O(n)$ algorithm from the previous section to answer the query "is there a set H with at most γ gaps and with $F_{\max}(H) \leq f$?" If the answer is "yes", we let $q = l$, otherwise we let $p = l + 1$. This will give us an algorithm with running time $O(n \log n)$.

Discrete case. We show that we can solve the scheduling variant in time $O(n \log n)$ as well. The solution is similar to the one for the discrete case, with two modifications. The first modification concerns the set Φ of candidate values for the maximum flow. We show that Φ can be still expressed as an $X + Y$ set. The second modification involves using Algorithm MinGapMaxFlow to answer decision queries in the binary search, instead of the algorithm for the continuous model.

We claim that, without loss of generality, we can make several simplifying assumptions about an optimal schedule. First, by a simple exchange argument, we can assume that the jobs are scheduled in order $1, 2, ..., n$. Two, we can assume that if two jobs are released at the same time then they are scheduled in the same block. (Otherwise, we can move the jobs in a later block to an earlier block.) Three, we can also assume that any block in an optimal schedule contains a job scheduled at its release time, for otherwise we can shift this block to the left.

Using the above assumptions, we reason as follows. If jobs $i, i+1, ..., i'$ are released at the same time then, among these jobs, job i' has the largest flow. Thus the maximum flow will be realized by a job with maximum index among the jobs that are released at the same time. Consider some block with jobs $a, a+1, ..., b$. Let $i' \in \{a, a+1, ..., b\}$ be the job with the largest flow time in this block and let $i, i+1, ..., i'$ be all jobs with $r_i = r_{i+1} = ... = r_{i'}$. We want to show that the flow time of job i' is of some specific simple form. Denote by t the time when i' is scheduled. If r_i is in this block, then the flow of i' is $t - r_{i'} = r_i + i' - i$. If r_i is before this block then let j be any job in this block executed at its release time r_j. By our assumption about job ordering, we have $r_j > t$. Then the flow time of i' can be expressed as $t - r_{i'} = r_j - (j - i') - r_{i'} = (r_j - j) - (r_{i'} - i')$.

Now, take X to be the set of all values $r_j - j$ for $j = 1, 2, ..., n$ and $Y = (-X) \cup \{1, 2, ..., n\}$. We can sort X and Y in time $O(n \log n)$. By the previous paragraph, we only need to search for the optimal flow value in $\Phi = X + Y$. Analogously to the continuous case, we perform binary search in Φ, using the $O(n)$-time algorithm from [6] for selection in $X + Y$ matrices and Algorithm MinGapMaxFlow as the decision oracle at each step. The running time will be $O(n \log n)$.

12 Final Comments

We studied several scheduling problems for unit-length jobs where the number of gaps in the schedule or their size is taken into consideration. For all problems we discuss we provided polynomial time algorithms, with running times ranging from $O(n \log n)$ to $O(n^7)$.

Many open problems remain. The most fundamental and intriguing question is whether the running time for minimizing the number of gaps for unit jobs (see Section 3) can be improved to below $O(n^4)$. Speeding up the algorithms in Sections 4, 5, 6, and 7 would also be of considerable interest.

There is a number of other variants of gap scheduling, even for unit jobs, that we have not addressed in our paper. Here are some examples:

- The problem of maximizing the minimum gap. This is somewhat similar to the problem we studied in Section 7, but we are not sure whether our method can be extended to this model. (We remark here that, according to our definition, the minimum gap size cannot be 0. For the purpose of maximizing the minimum gap, one can also consider an alternative model where "gaps" of size 0 are taken into account.)

- The tradeoff between throughput and gap size. Here, one can consider either the lower or upper bound on the gap size.

- The tradeoff between flow time (total or maximum) and gap sizes. For example, one may wish to minimize the total flow time with all gaps not exceeding a specified threshold.

- The problems of maximizing the number of gaps or minimizing the maximum gap, studied in Sections 6 and 7, were motivated by applications where the schedule for high-priority jobs needs to contain gaps where low-priority jobs can be inserted. A more accurate model for such applications would be to require that each block is of length at most b, for some given parameter b. Testing feasibility, with this requirement, can be achieved in high-degree polynomial time by extending the techniques from Sections 3-7, but it would be interesting to see whether more efficient solutions exist.

Another direction of research would be to study variants of gap scheduling for jobs of arbitrary length, for models with preemptive or non-preemptive jobs. The algorithm for minimizing the number of gaps, for example, can be extended to jobs of arbitrary length [3] if preemptions are allowed, although its running time increases from $O(n^4)$ to $O(n^5)$.

References

[1] Eric Angel, Evripidis Bampis, and Vincent Chau. Low complexity scheduling algorithm minimizing the energy for tasks with agreeable deadlines. In *10th Latin American Theoretical Informatics Symposium (LATIN'12)*, pages 13–24, 2012.

[2] Philippe Baptiste. Scheduling unit tasks to minimize the number of idle periods: a polynomial time algorithm for offline dynamic power management. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 364–367, 2006.

[3] Philippe Baptiste, Marek Chrobak, and Christoph Dürr. Polynomial time algorithms for minimum energy scheduling. In *Proceedings of the 15th Annual European Symposium on Algorithms (ESA)*, pages 136–150, 2007.

[4] Marek Chrobak, Uriel Feige, Mohammad Taghi Hajiaghayi, Sanjeev Khanna, Fei Li, and Seffi Naor. A greedy approximation algorithm for minimum-gap scheduling. In *Proc. 8th International Conference on Algorithms and Complexity (CIAC'13)*, pages 97–109, 2013.

[5] Erik D. Demaine, Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Amin S. Sayedi-Roshkhar, and Morteza Zadimoghaddam. Scheduling to minimize gaps and power consumption. In *Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 46–54, 2007.

[6] Andranik Mirzaian and Eshrat Arjomandi. Selection in X+Y and matrices with sorted rows and columns. *Inf. Process. Lett.*, 20(1):13–17, 1985.

[7] Wikipedia. Point coordination function. `http://en.wikipedia.org/wiki/Point_coordination_function`.

[8] Gerhard J. Woeginger. Monge strikes again: optimal placement of web proxies in the internet. *Oper. Res. Lett.*, 27(3):93–96, 2000.

Appendix A Correctness Proof for Algorithm MinGaps

In this section we prove correctness of Algorithm MinGaps. We divide the proof into two parts: feasibility and optimality. For feasibility, we show that for all k, a, b there is a schedule for $\mathcal{J}_{k,a,b}$ with $G_{k,a,b}$ non-final gaps. The optimality proof then will show that there is no schedule with fewer non-final gaps, and that any schedule having the same number of non-final gaps does not end after $\mathsf{S}_{k,a,b}$.

Feasibility. To justify feasibility, we need to show that there is a schedule for $\mathcal{J}_{k,a,b}$ in $[r_a+1, r_b-1]$ that has $\mathsf{G}_{k,a,b}$ non-final gaps. The proof is by simple induction. Let $Q_{k,a,b}$ denote such a schedule, and suppose we constructed such schedules $Q_{k',a',b'}$ for $k' < k$ and all a', b'.

We have cases, depending on how the value of $\mathsf{G}_{k,a,b}$ was realized. Suppose first that $\mathsf{G}_{k,a,b}$ is minimized by the top expression, and choose the c that realizes the minimum in this expression. For this c, we have $c < k$, $r_a < r_c < r_b$ and $\mathsf{S}_{k-1,a,c} \geq r_c - 2$. By induction, there exist schedules $Q_{k-1,a,c}$ and $Q_{k-1,c,b}$. By the stretch assumption above, we can assume that $Q_{k-1,a,c}$ schedules the last job no earlier than at $r_c - 2$. In fact, we can assume that it schedules this job at exactly $r_c - 2$ (by the compression lemma in [3], which basically says that we can compress a schedule without adding any gap to it). Then we construct $Q_{k,a,b}$ by scheduling jobs from $\mathcal{J}_{k-1,a,c}$ according to $Q_{k-1,a,c}$, scheduling k at $r_c - 1$, scheduling c at r_c, and scheduling all jobs in $\mathcal{J}_{k-1,c,b}$ according to $Q_{k-1,c,b}$. Then the number of non-final gaps in $Q_{k,a,b}$ will be equal to the sum of the number of non-final gaps in $Q_{k-1,a,c}$ and $Q_{k-1,c,b}$.

The other cases is when $\mathsf{G}_{k,a,b}$ is minimized by the bottom expression. By induction, there is a schedule $Q_{k-1,a,b}$ that realizes $\mathsf{G}_{k-1,a,b}$ and ends at $\mathsf{S}_{k-1,a,b}$. We have two sub-cases. If $\mathsf{S}_{k-1,a,b} < r_k - 1$, $Q_{k,a,b}$ will schedule all jobs in $\mathcal{J}_{k-1,a,b}$ according to $Q_{k-1,a,b}$ and schedule k at $\min(r_b - 1, d_k)$. Then $Q_{k,a,b}$ will have one more non-final gap than $Q_{k-1,a,b}$, which agrees with our formula.

In the other sub-case, $\mathsf{S}_{k-1,a,b} \geq r_k - 1$. Then the argument is as follows. If $\mathsf{S}_{k-1,a,b} < r_b - 1$, schedule $\mathcal{J}_{k-1,a,b}$ according to $Q_{k-1,a,b}$ and schedule k at $\mathsf{S}_{k-1,a,b} + 1$. (Note that k will meet its deadline, because $\mathsf{S}_{k-1,a,b} + 1 \leq d_{k-1} + 1 \leq d_k$.) Otherwise, we have $\mathsf{S}_{k-1,a,b} = r_b - 1$. Using the compression lemma in [3], we can then shift $Q_{k-1,a,b}$ leftward by one slot, so that it ends at $r_b - 2$ and schedule k at $r_b - 1$. This way, $Q_{k,a,b}$ will have the same number of non-final gaps as $Q_{k-1,a,b}$.

Optimality. We prove optimality by induction on k. Suppose that for $k' \leq k-1$ all values $\mathsf{G}_{k',a',b'}$ are optimal. We now consider $\mathsf{G}_{k,a,b}$. If $k \notin \mathcal{J}_{k,a,b}$ then the optimal schedule for $\mathcal{J}_{k,a,b}$ is also optimal for $\mathcal{J}_{k-1,a,b}$, so by setting $\mathsf{G}_{k,a,b} = \mathsf{G}_{k-1,a,b}$, the algorithm will compute the optimal value correctly.

So assume now that $k \in \mathcal{J}_{k,a,b}$. We start by analyzing the structure of an optimal schedule Q for $\mathcal{J}_{k,a,b}$. Without loss of generality, we can assume that Q has the ED property. Suppose that k is scheduled in Q at some time t. Suppose also that k is not the last job in Q. If k is scheduled as a first or last job in a block, then we can move k to later without increasing the number of gaps. Therefore we can assume that k is strictly inside a block, that is the slots $t-1$ and $t+1$ are busy.

By the ED property, no other jobs are pending at time t, therefore all jobs released before t are completed before t. This means that $t + 1 = r_c$, for some $c < k$. This also implies that $Q[r_a + 1, r_c - 2]$ is a schedule for $\mathcal{J}_{k-1,a,c}$, and that $Q[r_c + 1, r_b - 1]$ is a schedule for $\mathcal{J}_{k-1,c,b}$. (Here, $Q[x, y]$ is Q restricted to interval $[x, y]$.)

With the above in mind, we consider two cases.

Case 1: There is an optimal schedule for $\mathcal{J}_{k,a,b}$ with k scheduled as the last job. Let's then call Q_g the restriction of that schedule to the sub-instance $\mathcal{J}_{k-1,a,b}$ that we get by removing the job k.

• If there is a gap right before k in one such schedule (1) :
Q_g cannot be better than the one that we compute, and the total value of $\mathsf{G}_{k,a,b}$ is therefore reached by our below expression.
• If there is no gap right before k in any such schedule (2) :
Then we have $\mathsf{S}_{k-1,a,b}+1 \geq r_k$, for if it was not the case it would imply that Q_g has more gaps than $Q_{k-1,a,b}$ and adding the job k to the schedule $Q_{k-1,a,b}$ would be optimal for $\mathcal{J}_{k,a,b}$, which would contradict (2).

Case 2: All optimal schedules for $\mathcal{J}_{k,a,b}$ have k scheduled inside a block. So pick some optimal schedule Q. Let t be the time when k is scheduled. By the case assumption, k is scheduled inside a block in Q, $Q[r_a+1, r_c-2]$ is a schedule for sub-instance $\mathcal{J}_{k-1,a,c}$, and $Q[r_c+1, r_b-1]$ is a schedule for $\mathcal{J}_{k-1,c,b}$. The optimality of Q implies immediately that $Q[r_c+1, r_b-1]$ must be an optimal schedule for $\mathcal{J}_{k-1,c,b}$, so it has $\mathsf{G}_{k-1,c,b}$ gaps.

The argument for $Q[r_a+1, r_c-2]$ is more subtle. Here, the optimality of Q implies immediately that $Q[r_a+1, r_c-2]$ is an optimal schedule for $\mathcal{J}_{k-1,a,c}$ in $[r_a+1, r_c-1]$ with the restriction that the last jobs is scheduled at r_c-2. But this is not what we need; we need $Q[r_a+1, r_c-2]$ to actually be optimal for $\mathcal{J}_{k-1,a,c}$ in $[r_a+1, r_c-1]$, with no additional constraint. Let h be the number of gaps of $Q[r_a+1, r_c-2]$. Towards contradiction, supposed $Q[r_a+1, r_c-2]$ is not optimal. Then there is another schedule Q' for $\mathcal{J}_{k-1,a,c}$ in the interval $[r_a+1, r_c-1]$ with $h' < h$ gaps. This is not possible if Q' ends in r_c-2, by the earlier stated property of $Q[r_a+1, r_c-2]$. So Q' ends before r_c-2. However, in this case, we can move k to the end of Q, and as a result we would get a schedule whose number of non-final gaps does not exceed that in Q and has k scheduled at the end, which contradicts the assumption of the case (case 2).